

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Sorin Faibish, et al.

Serial No.: 09/893,825 Confirm: 5261

Filed: 06/28/2001

For: Video File Server Cache Management
 Using Movie Ratings for Reservation of
 Memory and Bandwidth Resources

Technology Center: 2600

Group Art Unit: 2623

Examiner: Parry, Christopher L.

Atty. Dkt. No.: 10830.0080.NPUS00

APPEAL BRIEF TO THE BOARD OF PATENT APPEALS AND INTERFERENCES

Commissioner for Patents
PO Box 1450
Alexandria, Virginia 22313-1450

Sir:

 This Appeal Brief is in support of the Notice of Appeal filed on Aug. 3, 2006. Please deduct any deficiency in any required fee from EMC Corporation Deposit Account No. 05-0889.

I. REAL PARTY IN INTEREST

The real party in interest is EMC Corporation, by virtue of an assignment recorded at Reel 011970 Frame 0829.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

III. STATUS OF THE CLAIMS

Claims 1-30 have been presented for examination.

Claims 1, 10, 15, 24, and 27-30 have been cancelled.

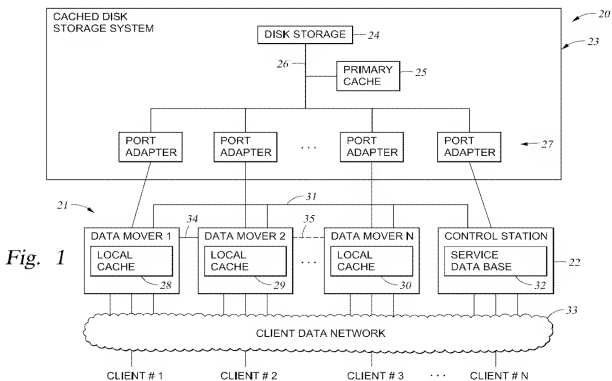
Claims 2-9, 11-14, 16-23, and 25-26 have been finally rejected, and are being appealed.

IV. STATUS OF AMENDMENTS

Appellants filed an Amendment After Final on July 10, 2006. This Amendment cancelled claims 27-30 and amended formal drawing FIG. 4 to show boxes of different sizes corresponding to the sizes shown in the informal drawings as originally filed. The Image File Wrapper in public PAIR includes an entry dated 07-18-2006 identified as “Amendment After Final or under 37 CFR 1.312 initialed by the examiner” and the initialed wording on the corresponding document (a copy of page 1 of appellants’ Amendment in Reply to Final Official Action) says: “Please Enter CLP 7/13/06”. Therefore, appellants understand that their Amendment After Final has been entered.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The appellants' invention is directed to a video file server (independent claims 2 and 12) and a method of operating the video file server (independent claims 16 and 26) for providing video-on-demand access to movies. The video file server includes a cached disk storage system including a primary cache and disk storage for storing the movies, and a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in the data network. Each of the data mover computers has a local cache. (Appellants' specification, page 4, lines 8-14.) This is shown in appellants' FIG. 1 as reproduced below and described in appellants' specification on page 7 line 12 et seq.:



The appellants' independent claims are more specifically directed to ranking the movies with respect to popularity, and pre-assigning a respective set of the data movers for servicing video streams for each movie ranking, wherein the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

The ranking of movies with respect to popularity is shown in appellants' FIG. 5, reproduced below:

MOVIE RANKING	MOVIE TITLE	FREQUENCY OF ACCESS	DATA MOVER SET	NO. OF RESERVED VIDEO STREAMS	NO. OF ACTIVE VIDEO STREAMS	LOCATION ON DISK	ANY LOCATION IN CACHE
1	—	25,000	1	5,000	3,497	—	—
2	—	11,000	2	2,200	1,738	—	—
3	—	7,900	3	1,580	2,300	—	—
4	—	6,900	3	1,380	790	—	—
5	—	4,000	4	800	570	—	—
6	—	2,600	4	520	640	—	—
7	—	2,300	4	460	198	—	—
8	—	2,000	4	400	340	—	—
9	—	1,700	5	340	270	—	—

Fig. 5

The process of ranking the movies with respect to popularity is described on page 11 lines 10-19 of appellants' original specification as follows:

When a movie is first released for VOD distribution, it will have an industry rating, such as the "Blockbuster" rating, with respect to other popular movies. Therefore, the newly released movie can be ranked with respect to the popular movies already stored

in the video file server, and the rank of each lower rated movie in the video file server can be decreased by one level as the newly released movie is written to disk storage of the video file server. New movie releases, for example, are written into the disk storage at the time of day of minimum demand, for example, about 3:00 a.m. Each day, the number of accesses for each movie is recorded to compute a running average of the frequency of access of each movie, in order to re-adjust the rank of each movie, and make effective the adjusted rank at the time of day of minimum demand.

The pre-assigning of a respective set of data movers for servicing video streams for each movie ranking is shown in appellants' FIG. 3, reproduced below:

MOVIE RANKING	DATA MOVER SET
1	SET 1 = {DM1, DM2, DM3}
2	SET 2 = {DM3, DM4}
3	SET 3 = {DM5}
4	SET 3 = {DM6}
5	SET 4 = {DM7}
6	SET 4 = {DM7}
7	SET 5 = {DM8}
8	SET 5 = {DM8}
⋮	⋮

Fig. 3

The pre-assigning of a respective set of data movers for servicing video streams for each movie ranking is described on page 11 line 20 to page 23 line 6 of appellants' original specification, as follows:

In accordance with one aspect of the present invention, for each movie ranking, a particular set of data mover resources are preassigned for servicing client access to the movie having the movie ranking. With reference to FIG. 3, for example, there is shown a table associating each movie ranking with a data mover set pre-assigned to service the movie having the associated ranking. Therefore, there is a certain maximum number of video streams that are available to each movie, corresponding to the total number of video streams that all of the data movers in the set can supply to the client data network. There is also a certain maximum amount of data mover local cache memory capacity available for servicing each movie, corresponding to the total amount of local cache memory in all of the data movers in the set.

The respective sets of data movers being configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies was shown in appellants' FIG. 4 as originally filed and as described on page 13 lines 15-23 of appellants' original specification, as follows:

As shown in FIG. 4, the physical configuration of each of the data movers 21 depends on the rank of the movie or movies to be serviced by the data mover. In particular, the data movers servicing the higher ranking (i.e., the most popular) movies have fewer cache RAM cards 41 and more network interface cards 42 than the data movers servicing the lower ranking (i.e., less popular) movies. This is a consequence of the fact that the cache RAM has sufficiently high bandwidth that a large number of video

streams can be serviced from one copy of a movie in the cache RAM. For less popular movies, not all of this bandwidth can be utilized, so that the ratio of network interface cards to cache RAM cards will fall for less popular movies.

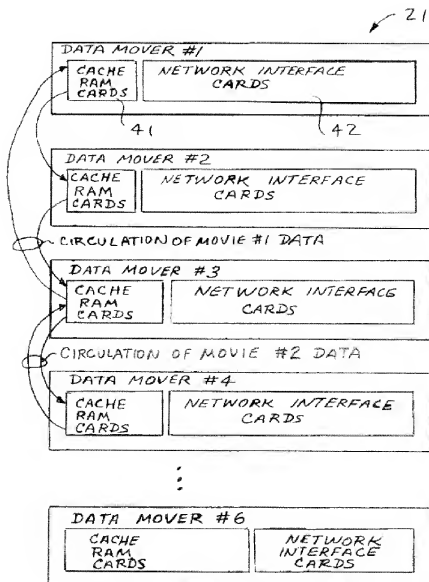


FIG. 4

Appellants' dependent claims 5 and 13 further define that a series of at least some of the data movers include direct links for transfer of movie data from a data mover set servicing one movie ranking to a data mover set servicing a next higher movie ranking and for transfer of movie data from the data mover set servicing the one movie ranking to a data mover set servicing a next lower movie ranking. Such direct links are links 34 and 35 shown in FIG. 1, reproduced above, and described in appellants' specification, page 13, lines 6-11:

For example, as shown in FIG. 1, the data movers are arranged in a series such that neighboring data movers in the series are coupled by a dedicated link 34, 35, etc. and the flow of movie data from one data mover to the next over a dedicated link permits the movie data from one data mover set servicing one movie ranking to be directly transferred to another data mover set servicing a next higher or lower movie ranking.

Appellants' independent claims 12 and 26 further recite "locking in the cache a plurality of entire movies, and when there is a need for servicing a more popular movie from the cache and there is insufficient free cache memory for servicing the more popular movie from the cache, transferring the servicing of a less popular movie from the cache to disk storage in order to free cache memory for servicing the more popular movie from the cache." (Appellants' specification, page 4, line 21, to page 5, line 3, and page 5, line 19, to page 6, line 2.) This is described in appellants' specification, page 22, lines 8-32, with reference to appellants' FIG. 20:

With reference to FIG. 20, in step 161, if the number of streams (tk) for the lowest ranking (i.e., least popular) movie in the primary cache is greater than a threshold, then execution continues to FIG. 14. In this case, it would be too

disruptive to move the origination of the streams from an entire movie in cache to the movie on disk. Otherwise, execution continues from step 161 to step 162. In step 162, execution branches to FIG. 14 if there is not enough bandwidth to remove the movie from locked primary cache and continue access of the movie from disk. This is done by comparing the sum of the bit rates for all of the streams from the lowest ranking (i.e., least popular) movie in the primary cache to the difference between the free disk bandwidth less the bit rate for the new stream. Otherwise, if there is enough bandwidth, execution continues to step 163. In step 163, the control station (CS) opens the file from disk for the movie M_k that was in the locked primary cache and is to be accessed from disk. In step 164, the control station (CS) sends a pointer to the data movers in the data mover set (DMs) servicing the movie M_k . In step 165, the movie M_k is removed from the primary cache. In step 166, the movie M_i is moved from disk to cache.

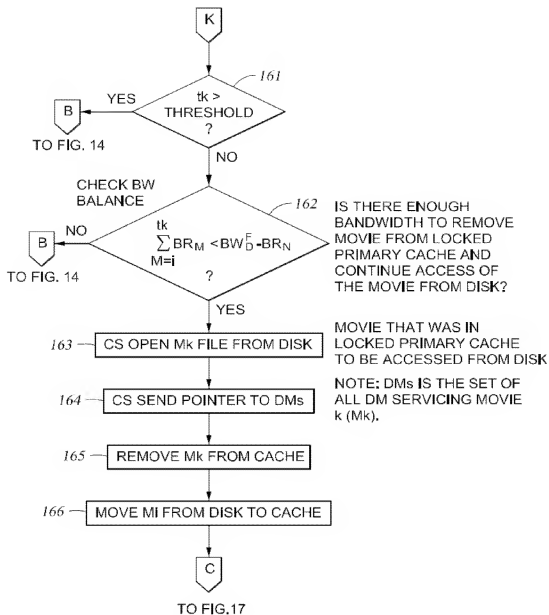


Fig. 20

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether claims 2-9, 11-14, 16-23, and 25-26 are unpatentable under 35 U.S.C. 103(a) over Armstrong et al. WO 2000/60861 in view of Mizutani U.S. Patent 6,115,750.

VII. ARGUMENT

Claims 2-9, 11-14, 16-23, and 25-26 are patentable under 35 U.S.C. 103(a) over Armstrong et al. WO 2000/60861 (hereinafter “Armstrong”) in view of Mizutani U.S. Patent 6,115,750.

The policy of the Patent and Trademark Office has been to follow in each and every case the standard of patentability enunciated by the Supreme Court in Graham v. John Deere Co., 148 U.S.P.Q. 459 (1966). M.P.E.P. § 2141. As stated by the Supreme Court:

Under § 103, the scope and content of the prior art are to be determined; differences between the prior art and the claims at issue are to be ascertained; and the level of ordinary skill in the pertinent art resolved. Against this background, the obviousness or nonobviousness of the subject matter is determined. Such secondary considerations as commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.

148 U.S.P.Q. at 467.

The problem that the inventor is trying to solve must be considered in determining whether or not the invention would have been obvious. The invention as a whole embraces the structure, properties and problems it solves. In re Wright, 848 F.2d 1216, 1219, 6 U.S.P.Q.2d 1959, 1961 (Fed. Cir. 1988).

Claims 2-4, 6-9, 11-12, 14, 16-18, 23, and 25-26

Armstrong discloses a method and apparatus for hierarchical distribution of video content for an interactive information distribution system. (Title.) A first embodiment shown in FIG. 1 comprises a remote server head-end 110R and a plurality of head-ends 110_i through 110_n each being coupled to at least one of a corresponding plurality of neighborhoods 130_i through 130_n. (Page 6, lines 13-19.) Each head-end comprises a host workstation 112, a video stream server 114, and a primary storage partition 118 comprised of an array of hard drives. (Page 6, lines 20-24.) Video streams are transmitted from the video stream server 114 to the subscriber's respective subscriber equipment comprised of a set-top box 142, a display 144 and a control device 145. (Page 6 line 32 to page 7 line 2.) The remote server head-end 110R also includes a secondary storage partition 119R coupled to the remote server 114R. (Page 7, lines 11-13.) A primary storage partition 118 of a head-end 110, including the remote server head-end 110R, is used to store frequently requested video assets. Alternately, the secondary storage partition 119 of the remote server head-end 110R is used to store infrequently requested video assets. (Page 7, lines 17-20.) The content manager 120 tracks the number of requests for a video asset and produces an asset request rate. An operator using the host workstation 112 defines a threshold rate for each video asset. The content manager 120 periodically compares the asset request rate against the threshold rate for each video asset in the system 100. If the asset request rate traverses the threshold rate for the video asset, then the video asset is stored on the primary storage partitions 118 and 118R at each of the head-ends 110 and 110R. If the asset request rate does not traverse the threshold rate for a video asset, then the video asset is stored on the secondary storage partition 119R at the remote server head-end 110R. In this

manner video assets are dynamically distributed throughout the interactive information distribution system 110. (Page 7, lines 21-31.)

Armstrong discloses a second embodiment in FIG. 2. In FIG. 2, the primary storage device 216 of the head-ends 210 is apportioned into at least two storage partitions designated as a primary storage partition 218, and a secondary storage partition 219. (Page 9, lines 14-16.) The primary storage partition 218 on the primary storage device 216 at each head-end 210 is used to store frequently requested video assets and temporarily cached library video assets. Each primary storage partition 218 at each head-end typically has the same frequently requested video assets as any other head-end 210. The secondary storage partition 219 is used to store portions of the infrequently requested video assets. An entire library of infrequently requested video assets is divided and stored amongst the plurality of head-ends 210 at each of the secondary storage partitions 219 on their respective primary storage devices 216. An infrequently requested video asset is typically stored on the secondary storage partition 219 at a single head-end 210. However, the request rate for that video asset may warrant additional storage at other head-ends. As such the content may be replicated and stored thereafter. In this manner, video assets that do not warrant storage across the entire system of head-ends 210 in the interactive information distribution system 200, may still be dynamically stored at multiple head-ends 210. Such dynamic storage corresponding to those neighborhoods having hither request rates than others is made in accordance with an algorithm that allows maximum access to the video titles with minimum network cost associated with their delivery. (Page 10, lines 1-19.) A threshold rate is a value for each requested video asset, established by the service provider in the interactive information distribution system 100, which

defines a level to be considered as frequent or infrequent request by the subscribers. Each video asset may have multiple threshold rates. Multiple threshold rates are set to establish various parameters for the storage locations of video information. Such parameters include discarding the video asset, storing it as a single head-end 110, replicating the video asset and storing it at more than one head-end 110 where the request rate warrants it, or storing it as all the head-ends 110 across the entire interactive information distribution system 100. (Page 14, lines 1 to 28.)

Page 8 of the final Official Action says: “Armstrong teaches, ‘wherein the movies are ranked with respect to popularity ...’ by disclosing primary storage partition 218 is used to store frequently requested video assets and secondary storage partition 219 is used to store infrequently requested video assets (page 10, lines 1-10).” Appellants respectfully disagree. Ranking is different from simply classifying movies as either frequently requested or not frequently requested. The plain meaning of the verb “rank” is “To arrange in a series in ascending or descending order of importance.” (See, for example, the definition on page 825 of Rudolf F. Graf, *Modern Dictionary of Electronics*, Butterworth-Heinemann, Newton, Ma 1997, a copy of which is attached hereto in the Evidence Appendix.) Such ranking of movies is shown in appellants’ FIG. 5.

In response to appellants’ argument, page 2 of the final Official Action refers to a definition of “rank” from the Webster’s New World Dictionary. This definition, however, is for the word “rank” used as a noun. The appellants’ independent claims are using the word “ranked” as a verb, and further specify a respective set of the data movers are assigned for

servicing video streams for each movie ranking, and the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies. The process defined in appellants' independent claims is not simply giving precedence or more resources to more popular movies than less popular movies. The process includes ranking the movies with respect to popularity, assigning a respective set of the data movers for servicing video streams for each movie ranking, and configuring the data movers in the respective sets of data movers differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

Page 8 of the final Official Action says: "Armstrong teaches, 'wherein the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies' by disclosing headend 210₂-210_n comprise primary storage partition 218 is used to store frequently requested video assets and secondary storage partition 219 is used to store infrequently requested video assets." Appellants respectfully disagree. It appears that each head-end data mover in Armstrong (FIGS. 1 and 2) has the same configuration with respect to cache resources and network interface resources. The appellants' data movers (FIG. 4) are configured differently by having fewer cache RAM cards and more network interface cards in the data movers assigned to storing and servicing the more popular movies than in the data movers assigned to storing and servicing the less popular movies. (Appellant's original specification, page 13 lines 15-

23.) In addition, the portions of Armstrong reproduced above suggest that more of the local cache memory will be used for storing more popular movies than less popular movies.

In the response to the appellants' argument, page 3 of the final Official Action notes that Armstrong teaches when a subscriber requests to view a frequently requested video asset or "popular movie" that is stored on primary storage partition 218, the movie is immediately delivered to the requesting subscriber. On the other hand, if a subscriber requests to view an infrequently requested video asset or "less popular movie" which is not stored on primary storage partition, 218, then resources for the less popular movie must be retrieved from other headends with the system. This is true if the less popular movie is not already in the secondary storage partition, but the primary storage partition 218 is not a network interface resource; instead, the primary storage partition 218 is a memory resources, as can be seen since the primary storage 216 is subdivided or partitioned into the primary storage partition 218 for the more popular video assets and the secondary storage partition 219 for the less popular video assets. In the context of appellants' specification (page 10, lines 10-11), the primary storage 216 of a head-end serves as a local cache for the head-end since the more popular movies are "kept" in the local cache by keeping them in the primary storage partition. Nor should a video asset or movie that would be retrieved from a network and stored in cache be considered to be a cache memory resource, since appellants' claims call for "local cache memory resources for less popular movies." More importantly, the appellants' independent claims are not specifying that one part of a data mover services popular movies and another part of the data mover services less popular movies. Instead the appellants' independent claims specify that data movers that service more popular movies have fewer cache

memory resources and more network interface resources than data movers that service less popular movies. (This is how the Examiner construed these claims on pages 6 to 7 of the final Official Action in his objection to claims 27-30 under 37 C.F.R. 1.75(c).)

Where the prior art references fail to teach a claim limitation, there must be “concrete evidence” in the record to support an obviousness rejection. “Basic knowledge” or “common sense” is insufficient. In re Zurko, 258 F.3d 1379, 1385-86, 59 U.S.P.Q.2d 1693, 1697 (Fed. Cir. 2001).

Page 8 of the final Official Action says: “Armstrong fails to disclose a respective set of data movers pre-assigned for servicing video streams for each movie ranking.” The final Official Action cites Mizutani for this feature.

Mizutani (US 6,115,740) discloses a video file server system for dynamically allocating contents and delivering data. (See title.) The video server system has a plurality of video servers having respective contents storing units for storing contents and respective contents delivering units for delivering contents. A management server has a stream supply information managing unit for managing stream supply information relative to the delivery of the contents and a contents dynamic allocating unit for controlling the storage of the contents between the video servers to dynamically allocate contents based on stream supply information from the stream supply information managing unit. (Abstract.) Mizutani says that different kinds of content C0, C1, C2, ... (e.g., different digitally moving image data, col. 1, lines 15-16) can be stored in a video server system. Each video server in the system can deliver a maximum number (Nstrm) of streams. (Col. 1, lines 48-49.) The maximum number of streams of content that can be

delivered at one time from an entire video server system may be increased by increasing the number of installed video servers. (Col. 1, lines 30-33.) In order to avoid rejection of a request for the delivery of a content stream C1, it is necessary that the content C1 be stored in the video server beforehand in expectation of access to the content C1. (Col. 1, lines 1-4.) The estimated number of video servers which can be installed (N_{vsa}) is the sum for all i of P_i/N_{strm} , where P_i represents the maximum number of times that each of the contents is simultaneously accessed per day, and i represents the type of a content. (Col. 2, lines 9-20.)

Mizutani says it has been customary to predict concentrated access to certain contents, estimate the number of video servers to be installed, and statistically allocate appropriate contents in the video servers before the video system is put into service. (Col. 3, lines 1-5.) Mizutani says that this conventional static contents allocation scheme usually results in an excessive estimate of the required number of servers to be installed (col. 3, lines 36-37) because of an incorrect assumption that the maximum numbers P_i of times that the respective contents C_i are simultaneously accessed occur at a common time (col. 3, lines 6-11). Instead, in normal circumstances, different users access different kinds of contents at different times. For example, news programs are popular in the morning, and movies are popular in the evening, and some users prefer to see video programs early in the evening and others late in the evening. (Column 3, lines 30-35.)

Mizutani's solution to the problems of the conventional static contents allocation scheme (col. 3, lines 26-28) is to dynamically allocate the contents (col. 3, lines 48-56). The contents are dynamically allocated by detecting whether at least the number of streams of a content stored in

a video server or the predictable number of accesses exceeds a corresponding threshold value or not, and if the number exceeds the threshold value, controlling the storage of the content between the video server or another video server, for thereby dynamically allocating the content. (Column 4, lines 3-11.) Predicted values used by the video file server system for dynamically allocating contents include a predicted maximum number $A(s,t)$ of times that a video server s is simultaneously accessed at a time t . (Col. 6, lines 6-24.) If there is a request from a user at time t , then the video server whose $A(s,t)$ is the smallest serves as a delivering video server for delivering a requested content. (Col. 6, lines 25-27.) Another predicted value is a number $B(i,t)$ of lacking resources of the content i predicted at the time t . (Col. 6, lines 28-29.) The predicted maximum number $B(i,t)$ is periodically checked for all contents, and the contents are dynamically allocated by being copied, moved, and deleted so that $B(i,t)=0$ as much as possible. (Col. 6, lines 51-54.) Contents are allocated according to a video server having smallest predicted number of simultaneous accesses at a given time. (Col. 14, lines 8-18.)

Pages 8 to 9 of the final Official Action says: “Mizutani teaches using the predicted number of times that the content I is simultaneously accessed at the time t is represented by $P_i(t)$ and the equation for $B(i,t)$ is used to determine if content is lacking resources to determine how many streams on each server are necessary to facilitate requests (Col. 6, lines 32-38).” As introduced above, however, $B(i,t)$ is a number of lacking resources of the content i predicted at the time t (Col. 6, lines 28-29), and ranking is different from a number of lacking resources.

Page 9 of the final Official Action says: “Figure 7 [of Mizutani] further discloses pre-assigning content, C_0 and C_1 , to servers SV_0 and SV_1 or ‘data movers’. Accordingly, it would

have been obvious to one of ordinary skill in the art at the time the invention was made to modify Armstrong with the teachings of Mizutani in order to preassign data movers to service video streams for the benefit of making more resources available for more popular content.” However, the appellants’ invention of claim 2 does not result merely by assigning copies of popular content to be stored in the local cache of particular server head-ends in FIG. 1 or FIG. 2 of Armstrong. Nor does the prior art as a whole suggest that Armstrong and Mizutani should be combined or modified as required to arrive at the appellants’ invention.

As introduced above, Armstrong uses a threshold technique for hierarchical distribution of video content in a video-on-demand system. Armstrong’s technique appears entirely suitable for its intended purpose.

Mizutani finds fault with the prior art static contents allocation scheme (col. 3, lines 26-28) and teaches instead a method of dynamically allocating contents. It is respectfully submitted that from the viewpoint of FIG. 16 as a point of origin, Mizutani and the appellants go off in different directions in an attempt to provide more efficient allocation of video server resources and thus avoid installation of an excessive number of video servers to satisfy client demand. Mizutani does not appear to care which contents are popular and which are not, because the invention of Mizutani should dynamically move content between the servers to suit changing conditions. A reference such as Mizutani should be considered as a whole, and portions arguing against or teaching away from the claimed invention must be considered. Basch & Lomb, Inc. v. Barnes-Hind/Hydrocurve, Inc., 796 F.2d 443, 230 U.S.P.Q. 416 (Fed. Cir. 1986), cert. denied, 484 U.S. 823 (1987).

In short, neither Armstrong nor Mizutani suggests that the content should be ranked, a respective set of servers should be pre-assigned for servicing video streams for each movie ranking, and the servers in the respective sets of servers should be configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

Hindsight reconstruction, using the applicant's specification itself as a guide, is improper because it fails to consider the subject matter of the invention "as a whole" and fails to consider the invention as of the date at which the invention was made. "[T]here must be some motivation, suggestion, or teaching of the desirability of making the specific combination that was made by the applicant." In re Lee, 277 F.3d 1338, 1343, 61 U.S.P.Q.2d 1430, 1435 (Fed. Cir. 2002) (quoting In re Dance, 160 F.3d 1339, 1343, 48 U.S.P.Q.2d 1635, 1637 (Fed. Cir. 1998)). "[T]eachings of references can be combined only if there is some suggestion or incentive to do so." In re Fine, 837 F.2d 1071, 1075, 5 U.S.P.Q.2d 1596, 1600 (Fed. Cir. 1988) (Emphasis in original) (quoting ACS Hosp. Sys., Inc. v. Montefiore Hosp., 732 F.2d 1572, 1577, 221 U.S.P.Q. 929, 933 (Fed. Cir. 1984)). "[P]articular findings must be made as to the reason the skilled artisan, with no knowledge of the claimed invention, would have selected these components for combination in the manner claimed." In re Kotzab, 217 F.3d 1365, 1371, 55 U.S.P.Q.2d 1313, 1317 (Fed. Cir. 2000). See, for example, Fromson v. Advance Offset Plate, Inc., 755 F.2d 1549, 1556, 225 U.S.P.Q. 26, 31 (Fed. Cir. 1985) (nothing of record plainly indicated that it would have been obvious to combine previously separate lithography steps into one process); In re Gordon et al., 733 F.2d 900, 902, 221 U.S.P.Q. 1125, 1127 (Fed. Cir. 1984) (mere fact that prior

art could be modified by turning apparatus upside down does not make modification obvious unless prior art suggests desirability of modification); Ex Parte Kaiser, 194 U.S.P.Q. 47, 48 (PTO Bd. of Appeals 1975) (Examiner's failure to indicate anywhere in the record his reason for finding alteration of reference to be obvious militates against rejection).

Claims 5 and 13

Dependent claims 5 and 13 further define that a series of at least some of the data movers include direct links for transfer of movie data from a data mover set servicing one movie ranking to a data mover set servicing a next higher movie ranking and for transfer of movie data from the data mover set servicing the one movie ranking to a data mover set servicing a next lower movie ranking. Page 11 of the final Official Action says: "Armstrong fails to explicitly disclose transferring the movie data to a data mover servicing a next higher/lower movie ranking." It is also not clear whether the links in Armstrong between the head end servers are direct or not. Appellants also disagree with the contention that it is notoriously well known in the art to transfer movie data to servers serving a next higher/lower movie ranking, since the examination should be based on actual evidence, and in the context of appellants' specification, ranking is more than simply classifying movies as either frequently requested or not frequently requested. Moreover, page 11 of the Final Official Action refers to a "benefit of having servers with more or less resources service movies that need more or less resources depending on popularity." However, neither Armstrong nor Mizutani discloses servers that are configured differently for providing one set of servers that are better suited for serving more popular movies and another

set of servers that are better suited for serving less popular movies. Therefore it is respectfully submitted that the Final Official Action has not provided a sufficient basis for concluding that it would have been obvious to modify Armstrong to include the features of appellants' claims 5 and 13 that are admittedly missing from Armstrong.

In view of the above, the rejection of claims 2-9, 11-14, 16-23, and 25-26 should be reversed.

Respectfully submitted,

A handwritten signature in dark ink, appearing to read "Richard C. Auchterlonie", is written over a light gray circular stamp.

Richard C. Auchterlonie
Reg. No. 30,607

NOVAK DRUCE & QUIGG, LLP
1000 Louisiana, 53rd Floor
Houston, TX 77002
713-571-0655

VIII. CLAIMS APPENDIX

The claims involved in this appeal are as follows:

2. A video file server for providing clients with video-on-demand access to movies, the video file server comprising:

a cached disk storage system including a primary cache and disk storage for storing the movies; and

a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in a data network, each of the data mover computers having a local cache;

wherein the movies are ranked with respect to popularity, and a respective set of the data movers are pre-assigned for servicing video streams for each movie ranking; and

wherein the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

3. The video file server as claimed in claim 2, wherein for very popular movies, the very popular movies are retained in their entirety in local cache of the data movers assigned to service the very popular movies.

4. The video file server as claimed in claim 2, wherein the sets of data movers include a set consisting of more than one data mover for servicing one very popular movie, a set consisting of one data mover for servicing only one movie, and a set consisting of one data mover for servicing a plurality of the movies.

5. The video file server as claimed in claim 2, wherein a series of at least some of the data movers include direct links for transfer of movie data from a data mover set servicing one movie ranking to a data mover set servicing a next higher movie ranking and for transfer of movie data from the data mover set servicing the one movie ranking to a data mover set servicing a next lower movie ranking.

6. The video file server as claimed in claim 2, wherein data mover resources for a certain number of video streams from the data movers to the clients are reserved for each of a multiplicity of the movies.

7. The video file server as claimed in claim 2, wherein the video file server is programmed for locking in the primary cache a plurality of entire movies, and when there is a need for servicing a more popular movie from the primary cache and there is insufficient free cache memory for servicing the more popular movie from the primary cache, transferring the servicing of a less popular movie from the primary cache to disk storage in order to free cache memory for servicing the more popular movie from the primary cache.

8. The video file server as claimed in claim 7, wherein the video file server is programmed for freeing primary cache memory by transferring the servicing of a least popular movie in the primary cache from the primary cache to the disk storage so long as no more than a certain number of video streams are being serviced concurrently from the least popular movie in the primary cache.

9. The video file server as claimed in claim 2, wherein the video file server is programmed for negotiating with a client for selection of an available movie during peak demand when resources are not available to select freely any movie in the disk storage for which a video stream can be started.

11. The video file server as claimed in claim 12, wherein the video file server is programmed for freeing locked cache memory by transferring the servicing of the least popular movie in the cache from the cache to the disk storage so long as no more than a certain number of video streams are being concurrently serviced from the least popular movie in the cache.

12. A video file server for providing clients with video-on-demand access to movies, the video file server comprising:

a cached disk storage system including a cache and disk storage for storing the movies;
and

a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in a data network;

wherein the video file server is programmed for locking in the cache a plurality of entire movies, and when there is a need for servicing a more popular movie from the cache and there is insufficient free cache memory for servicing the more popular movie from the cache, transferring the servicing of a less popular movie from the cache to disk storage in order to free cache memory for servicing the more popular movie from the cache,

wherein each of the data mover computers has a local cache, the movies are ranked with respect to popularity, and a respective set of the data movers are pre-assigned for servicing video streams for each movie ranking, and the data movers in the respective sets of data movers are configured differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

13. The video file server as claimed in claim 12, wherein a series of at least some of the data movers include direct dedicated links for transfer of movie data from a data mover set servicing one movie ranking to a data mover set servicing a next higher movie ranking and for transfer of movie data from the data mover set servicing the one movie ranking to the data mover set servicing a next lower movie ranking.

14. The video file server as claimed in claim 12, wherein data mover resources for a certain number of video streams from the data movers to the clients are reserved for each of a multiplicity of the movies.

16. A method of operating a video file server for providing clients with video-on-demand access to movies, the video file server having a cached disk storage system including a primary cache and disk storage containing the movies, and a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in a data network, each of the data mover computers having a local cache, wherein the method includes:

ranking the movies with respect to popularity, and assigning a respective set of the data movers to each movie ranking, and

servicing video streams for each movie ranking with the respective set of data movers assigned for servicing said video streams for said each movie ranking; and

which includes configuring differently the data movers in the respective sets of data movers in order to provide more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

17. The method as claimed in claim 16, which includes, for very popular movies, retaining the very popular movies in their entirety in the local cache of the data movers assigned to service the very popular movies.

18. The method as claimed in claim 16, which includes servicing a most popular movie with an assigned data mover set consisting of more than one data mover, servicing only one movie with an assigned data mover set consisting of one data mover, and servicing a plurality of movies with an assigned data mover set consisting of one data mover.

19. The method as claimed in claim 16, wherein a series of at least some of the data movers are linked by direct dedicated data links and the method includes transferring movie data from a data mover set servicing one movie ranking to a data mover set servicing a next higher movie ranking and transferring movie data from a data mover set servicing the one movie ranking to a data mover set servicing a next lower movie ranking.

20. The method as claimed in claim 16, which includes reserving data mover resources for a respective number of video streams from the data movers to the clients for each of a multiplicity of the movies.

21. The method as claimed in claim 16, which includes locking in the primary cache a plurality of entire movies, and when there is a need for servicing a more popular movie from the primary cache and there is insufficient free cache memory for servicing the more popular movie from the primary cache, transferring the servicing of a less popular movie from the primary cache to the disk storage in order to free primary cache memory for servicing the more popular movie from the primary cache.

22. The method as claimed in claim 16, which includes freeing primary cache memory by transferring the servicing of a least popular movie in the primary cache from the primary cache to the disk storage so long as no more than a certain number of video streams are being concurrently serviced from the least popular movie in the primary cache.

23. The method as claimed in claim 16, which includes the video file server negotiating with a client for selection of an available movie during peak demand when resources are not available to select freely any movie in the disk storage for which a video stream can be started.

25. The method as claimed in claim 26, which includes the video file server freeing locked cache memory by transferring the servicing of a least popular movie in the cache from the cache to the disk storage so long as no more than a certain number of video streams are being concurrently serviced from the least popular movie in the cache.

26. A method of operating a video file server for providing clients with video-on-demand access to movies, the video file server having a cached disk storage system including a cache and disk storage containing the movies, and a multiplicity of data mover computers coupled to the cached disk storage system for streaming video data from the cached disk storage system to clients in a client data network, the method comprising:

locking in the cache a plurality of entire movies, and when there is a need for servicing a more popular movie from the cache and there is insufficient free cache memory for servicing the more popular movie from the cache, transferring the servicing of a less popular movie from the cache to the disk storage in order to free cache memory for servicing the more popular movie from the cache,

wherein each of the data mover computers has a local cache, the method includes ranking the movies with respect to popularity, assigning a respective set of the data movers for servicing

video streams for each movie ranking, and configuring the data movers in the respective sets of data movers differently for providing more network interface resources for very popular movies and for providing more local cache memory resources for less popular movies.

IX. EVIDENCE APPENDIX

Following is a copy of page 825 of Rudolf F. Graf, Modern Dictionary of Electronics, as entered into the record on Feb. 2, 2006 with the Remarks/Argument in Appellants' Amendment in Reply to Official Action.

MODERN
DICTIONARY
of
ELECTRONICS

SIXTH EDITION
REVISED AND UPDATED

Rudolf F. Graf

Newnes
Boston Oxford Johannesburg Melbourne New Delhi Singapore

Newnes is an imprint of Butterworth-Heinemann

Copyright © 1997 by Butterworth-Heinemann

 A member of the Reed Elsevier group

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

- Recognizing the importance of preserving what has been written, Butterworth-Heinemann prints its books on acid-free paper whenever possible.

Library of Congress Cataloging-in-Publication Data

Graf, Rudolf F.

Modern dictionary of electronics / by Rudolf F. Graf. — 6th ed.
p. cm.

Reprint. Originally published: Indianapolis, Ind. : H W Sams,

© 1984.

ISBN 0-7506-9870-5

1. Electronics—Dictionaries.

TK7804 G67 1996

96-25894

621.381'03—dc20

CIP

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

The publisher offers special discounts on bulk orders of this book.

For information, please contact:

Manager of Special Sales

Butterworth-Heinemann

313 Washington Street

Newton, MA 02158-1626

Tel: 617-928-2500

Fax: 617-928-2620

For information on all Newnes electronics publications available, contact our World Wide Web home page at: <http://www.bh.com/bh>

Edited by: *Charlie Buffington and Jack Davis*

Illustrated by: *T.R. Enrick*

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

We are in the midst of a high-advances in electronics and closely related fields. To keep up with this exciting industry, we must not only record additional terms but also explain them with authority in a way that is quickly suited to its time and place.

Those who work in new areas must effectively communicate thoughts and information. Originators of newly coined definitions frequently change with actual use by others.

Every new edition of this dictionary of the electronics industry. It is a dictionary dedicated to the task of communication, deriving its authority from a clear and simple style that is understandable of complexity of the term being defined. This sixth edition of the dictionary probably the most up-to-date electronic definitions of approximate and related fields. This includes 5000 terms from the fifth edition published in 1977, as well as the first edition published online. The first edition published online were reviewed and revised or enhanced the intelligibility of each definition requiring no further updates, modified and augmented.

While this work is as up-to-date as the field of electronics is expanding, it evolves and established terms take on new meanings. The publisher intends to issue periodic updates; thus suggestions for new terms are welcomed.

824

random-number generator—A special machine routine or hardware that produces a random number or series of random numbers in accordance with specified limitations.

random pard—Pertains to that portion of the total pard in an electric power supply which is not periodic. This phenomenon is frequently referred to as noise.

random processing—The treatment of information without respect to where it is located in external storage and in an arbitrary sequence determined by the input against which it is to be processed.

random pulsing—Varying the repetition rate of pulses by noise modulation or continuous frequency change.

random sample—A sample in which every item in the lot is equally likely to be selected in the sample.

random sampling—1. A sampling process in which there is a significant time uncertainty between the signal being sampled and the taking of samples. 2. A selection of observations taken from all of the observations of a phenomenon in such a way that each chosen observation has the same possibility of selection.

random-sampling oscilloscope—An oscilloscope that functions by constructing a coherent display from samples taken at random.

random sequential memory—A memory in which one reference can be found immediately; the other reference is found in a fixed sequence.

random signals—Waveforms having at least one parameter (usually amplitude) that is a random function of time (e.g., thermal noise or shot noise).

random variable—1. Also called *variate*. The result of a random experiment. 2. A discrete or continuous variable which may assume any one of a number of values, each having the same probability of occurrence. 3. Also called *stochastic variable*. Any signal the amplitude or phase of which cannot be predicted by a study of previous values of the signal.

random variation—A fluctuation in data which is due to uncertain or random occurrences.

random velocity—The instantaneous velocity of a particle without regard to direction. It may be characterized by its distribution function or by its average, root-mean-square, or most probable value.

random vibration—A vibration generally composed of a broad, continuous spectrum of frequencies, the instantaneous magnitude of which cannot be specified at any given moment of time. Instantaneous amplitude can only be defined statistically by a probability distribution function that gives the fraction of the total time that the amplitude lies within

825

specified amplitude intervals. (If random vibration has instantaneous magnitudes distributed according to the Gaussian distribution, it is called *Gaussian random vibration*.) See also *White Noise*.

random winding—A coil winding in which the turns and layers are not regularly positioned or spaced but are positioned haphazardly.

random wound—Describing a coil wound without care to ensure that the wire is in layers. Random-wound coils have fewer turns for a given volume.

range—1. The maximum useful distance of a radar or radio transmitter. 2. The difference between the maximum and the minimum value of a variable. 3. The set of values that may be assumed by a quantity or function. 4. See *Receiving Margin*.

range-amplitude display—A radar display in which a time base provides the range scale from which echoes appear as deflections normal to the base.

range calibration—Adjustment of radar-range indications by use of known range targets or delayed signals so that, when on target, the radar set will indicate the correct range.

range coding—A method of coding a beacon response so that the response appears as a series of pulses on a radar scope. The coding provides identification.

range finder—1. A movable, calibrated unit of the receiving mechanism of a teleprinter that can be used to move the selecting interval relative to the start signal. 2. An optical distance finder that depends on triangulation of two convergent beams on an object from disparate viewpoints. A pair of unaided human eyes coupled with a computerized brain can estimate distance, at least for nearby objects, with some accuracy. 3. A device that depends on the measurement of time of wave travel from an object to a point, as in radar and sonar.

range gate—A gate voltage used to select radar echoes from a very short-range interval.

range-height indicator—A radar display on which an echo appears as a bright spot on a rectangular field. The slant range is indicated along the x-axis, and the height above the horizontal plane (on a magnified scale) along the y-axis. A cursor shows the height above the earth. Abbreviated *rh*.

range mark—See *Distance Mark*.

range marker—A variable or movable discontinuity in the range time base of a radar display (in the case of a ppi, a ring). It is used for measuring the range of an echo or calibrating the range scale.

range of an instrument—See *Total Range of an Instrument*.

range of gain—The minimum and maxi-

random winding—raster display

imum gain to which the amplifier can be set.

range resolution—The minimum difference in range between two radar targets along the same line of bearing for which an operator can distinguish between targets.

range ring—An accurate, adjustable ranging mark on a plan-position indicator corresponding to a range set on a type-M indicator.

range step—The vertical displacement on an M-indicator sweep to measure range.

range surveillance—Surveillance of a missile range by means of electronic and other equipment.

range unit—A radar-system component used for control and indication (usually counters) of range measurements.

range zero—Alignment of the start of a sweep trace with zero range.

ranging oscillator—An oscillator circuit containing an LC resonant combination in the cathode circuit, usually used in radar equipment to provide range marks.

rank—To arrange in a series in ascending or descending order of importance.

rapid memory—See *Rapid Storage*.

rapid storage—Computer storage in which the access time is very short; rapid access usually is gained by limiting the storage capacity. Also called *rapid memory*, *fast-access storage*, and *high-speed storage*.

rare gas—See *Noble Gas*.

raster—Acronym for Radio Amplification by Stimulated Emission of Radiation, a chemical "pumping" process that is accomplished without external radiation.

raster—1. On the screen of a cathode-ray tube, a predetermined pattern of scanning lines which provide substantially uniform coverage of an area. 2. The illuminated area produced by the scanning lines on a television picture tube when no signal is being received. 3. Rectangular line pattern of light produced on the screen of a cathode-ray tube with no signal present. It is formed by deflecting the electron beam rapidly from left to right and relatively slowly from top to bottom. 5. The pattern of lines traced by rectilinear scanning in display systems.

raster burn—In camera tubes, a change in the characteristics of the area that has been scanned. As a result, a spurious signal corresponding to that area will be produced when a larger or tilted raster is scanned.

raster display—1. A display in which the entire display surface is scanned at a constant refresh rate. 2. A refresh graphics system where the electron beam sweeps horizontally across the face of the crt from left to right, drawing the picture as a series of scan lines. At the end of each line, the beam is turned off and repositioned.

X. RELATED PROCEEDINGS APPENDIX

None.